



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Structural Attention: Rethinking Transformer for Unpaired Medical Image Synthesis^{*}

Vu Minh Hieu Phan¹, Yutong Xie¹, Bowen Zhang¹, Yuankai Qi², Zhibin Liao¹, Antonios Perperidis¹, Son Lam Phung³, Johan W. Verjans¹, and Minh-Son To⁴

¹ Australian Institute for Machine Learning, University of Adelaide, Australia

² Macquarie University, Australia

³ University of Wollongong, Australia

⁴ Flinders University, Australia

Abstract. Unpaired medical image synthesis aims to provide complementary information for an accurate clinical diagnostics, and address challenges in obtaining aligned multi-modal medical scans. Transformer-based models excel in imaging translation tasks thanks to their ability to capture long-range dependencies. Although effective in supervised training, their performance falters in unpaired image synthesis, particularly in synthesizing structural details. This paper empirically demonstrates that, lacking strong inductive biases, Transformer can converge to non-optimal solutions in the absence of paired data. To address this, we introduce UNet Structured Transformer (UNest) — a novel architecture incorporating structural inductive biases for unpaired medical image synthesis. We leverage the foundational Segment-Anything Model to precisely extract the foreground structure and perform structural attention within the main anatomy. This guides the model to learn key anatomical regions, thus improving structural synthesis under the lack of supervision in unpaired training. Evaluated on two public datasets, spanning three modalities, *i.e.*, MR, CT, and PET, UNest improves recent methods by up to 19.30% across six medical image synthesis tasks. Our code is released at <https://github.com/HieuPhan33/MICCAI2024-UNest>.

1 Introduction

Multi-modal medical imaging, including modalities such as computed tomography (CT), magnetic resonance imaging (MRI) and positron emission tomography (PET), can be a valuable tool with multiple applications in clinical practice, aiding in accurate disease diagnosis [4,7], lesion detection [5] and treatment planning [1]. While these modalities provide complementary insights, obtaining multiple scans can be time-consuming, costly, and, most importantly, potentially harmful to patients through additional radiation exposure [21] (*i.e.*, CT and PET). Medical image synthesis has emerged as a potential solution to enable comprehensive patient evaluation without requiring multiple scans.

^{*} Acknowledgement: This study was supported by the Channel 7 Children’s Research Foundation and the Royal Australian and New Zealand College of Radiologists.

Most synthesis methods employ a supervised Pix2Pix [9,6,31], showing effectiveness given paired data between two domains. Yet, acquiring substantial paired data is challenging due to variations in resolution and patient positioning. This highlights the advantage of unpaired image synthesis. CycleGAN [33] is a pioneering work in unpaired image translation, which imposes domain cycle consistency using two generators. Subsequent approaches further incorporate structural consistency loss, either by aligning pixel-wise correlations [27,28,17] or preserving shape fidelity [23,11] between original and synthesized images. Previous methods adopt convolution operators with the local inductive biases, guiding the model to extract local features. This limits their ability to capture long-range spatial contexts.




Fig. 1. (a) Synthetic MR-to-CT results of different ViT methods: ResViT [6], UNETR [12], and our proposed UNest, where our UNest most accurately preserves the structural cavity. (b) Attention maps for two patches in a smooth brain region (highlighted by a purple star) a structural nasal cavity (indicated by a red star). Transformer methods tend to focus on less relevant background features.

Recently, Vision Transformer (ViT) [8] discards the inductive biases to learn features globally. It has achieved state-of-the-art results for many supervised image-to-image translation tasks such as medical semantic segmentation [12,24], and image super-resolution [3,29]. ResViT [6] has advanced the use of hybrid Convolution-Transformer models for paired image translation. While several unpaired methods [22,25] develop hybrid models, their architectures primarily rely on convolutional operators.

Unfortunately, ViT models [12,6] struggle when applied to unpaired medical image synthesis. Fig. 1(a) shows examples that current ViT methods fail to synthesize intricate anatomical structures in the nasal cavities. Without convolution-like inductive biases, recent studies [16,20] show that ViT is less sample-efficient and fail to attend to discriminative features in low data regime. In this paper, we empirically discover that the ViT’s performance is non-optimal in unpaired image translation scenarios, where the lack of supervision signals exaggerates the problem. Fig. 1(b) shows that UNETR and ResViT attend to less relevant background regions when learning foreground structures in the brain and nasal

cavity (respectively denoted by the pink and red stars), suggesting non-optimal solution under unpaired training. Solutions such as Swin UNETR [24] use local attention on specific windows; yet, confining the context within a restricted window can exclude crucial anatomical regions.

This paper introduces a strong inductive bias into the Transformer and proposes a simple-yet-effective architecture, called UNet Structural Transformer (UNest) for effective unpaired medical image synthesis. At the core of UNest, our Structural Transformer (ST) block segregates foreground and background tokens, and injects structural and local inductive biases respectively on each region. When synthesizing foreground, our structural attention aggregates contexts within the anatomy to encode relationships between anatomical regions. Fig. 1(b) shows that UNest adaptively focuses on local areas when synthesizing structural cavities (pink) and more globally when generating smooth brain (red).

Our contributions can be summarized as follows. **1)** We show empirically that injecting a structural inductive bias enables Transformer to focus on discriminative areas, thus enhancing the synthesis of anatomical structures in unpaired image synthesis. **2)** We introduce a simple-yet-effective architecture, coined UNet Structural Transformer (UNest), applying a dual attention strategy: structural attention for the foreground and local attention for the background. **3)** Evaluated across six translation tasks covering three modalities: MR, CT, and PET, our method significantly improves the accuracy across various anatomical structures, from head to whole-body images.

2 Method

CycleGAN’s overview. Let $\{\mathbf{x}_n\}_{n=1}^n$ and $\{\mathbf{y}_m\}_{m=1}^M$ denote unpaired training samples from domains X and Y . Based on CycleGAN [33], our framework consists of two generators $G_{XY} : X \mapsto Y$ and $G_{YX} : Y \mapsto X$, learning forward and backward mappings between the two domains, as shown in Fig. 2(c). Here, G_{XY} and G_{YX} are trained to fool the respective discriminator D_Y and D_X via adversarial loss

$$\mathcal{L}_{\text{adv},Y}(G_X, D_Y) = \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} [\log D_Y(\mathbf{y})] + \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log(1 - D_Y(G_X(\mathbf{x})))] \quad (1)$$

Adversarial loss $\mathcal{L}_{\text{adv},X}$ for training generator G_Y is defined similarly. For unpaired training, CycleGAN imposes the cycle consistency loss

$$\mathcal{L}_{\text{cycle}} = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \|\mathbf{x} - G_Y(G_X(\mathbf{x}))\| + \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} \|\mathbf{y} - G_X(G_Y(\mathbf{y}))\|. \quad (2)$$

Analysis of Transformer models on unpaired image synthesis. Given a feature map $F \in \mathbb{R}^{H \times W \times K}$, the attention mechanism in Transformer [26] first projects a feature map F into query Q , key K , and value V . Given a query token $\mathbf{q} \in \mathbb{R}^K$, it aggregates the contextual information from a set of T surrounding key-value tokens $K_S \in \mathbb{R}^{T \times K}$, and $V_S \in \mathbb{R}^{T \times K}$. The global attention methods (e.g., UNETR [12], ResViT [6]) consider aggregating the entire image with a global scope S :

$$S = \{(i, j) | i \in [1, H], j \in [1, W]\}. \quad (3)$$

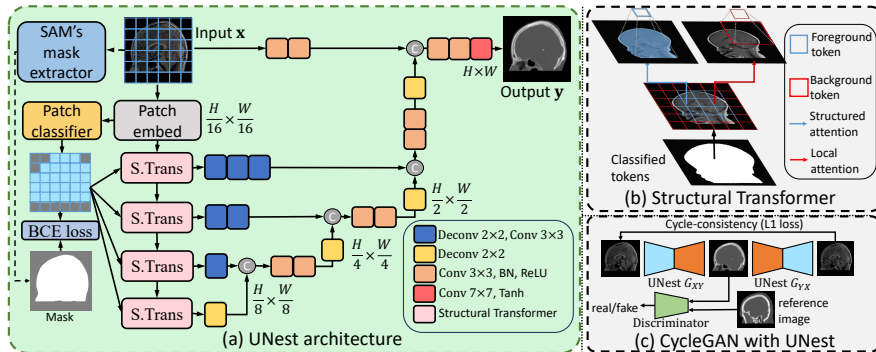


Fig. 2. (a) UNest architecture uses Structural Transformer blocks (pink block), as shown in (b), to perform dual-attention strategy on foreground and background tokens separately. The decoder upsamples features using deconvolutional layers and skip-connections from early encoder layers. (c) CycleGAN with UNest generators.

Learning on a global scope, Transformer is considered lacking local inductive bias. In contrast, the local attention methods (e.g., Swin UNETR) incorporate a local inductive bias, inspired by CNN. They aggregate tokens from a local $m \times m$ window surrounding the query token:

$$S = \{(i, j) | i \in [i_q - \frac{m}{2}, i_q + \frac{m}{2}], j \in [j_q - \frac{m}{2}, j_q + \frac{m}{2}]\}, \quad (4)$$

where (i_q, j_q) is the index of the query token q in the image. The self-attention module $SA(\cdot, \cdot)$ learns a new query feature by aggregating the key K_S and value vectors V_S within the pre-defined scope S as $\tilde{F} = SA(F, S) = \text{Softmax}(QK_S^T)V_S$.

Despite achieving high accuracy on supervised learning tasks [6,32,30,12], current Transformer models falter under unpaired image synthesis. Fig. 3 shows that global attention adopted in ResViT [6] and UNETR [12] deforms the hip structures, while the local attention in Swin UNETR [24] generates artifacts. This indicates that lacking strong inductive biases makes current Transformer models converge to a non-optimal solution under unpaired training.




Fig. 3. Visual results for PET-to-CT on AutoPET dataset [10].

Our UNet Structural Transformer. Given the lack of supervision in unpaired image translation, we propose to inject the structural inductive bias, enforcing the attention on discriminative anatomical structures. We introduce a UNest architecture, for unpaired medical image synthesis. UNest consists of Structural Transformer (ST) blocks and a convolutional decoder with skip connections, as shown in Fig. 2(a).

As shown in Fig. 2(a), Transformer-based encoder $E(\cdot)$ first splits the input image \mathbf{x} of shape $H \times W$ into a sequence of HW/p^2 non-overlapping patches with patch size $p \times p$. We flatten this sequence and denote it as $X_p \in \mathbb{R}^{(H' \cdot W') \times P}$, where $H' = \frac{H}{p}$, $W' = \frac{W}{p}$, and $P = p^2$ is the total size of each patch token. Each patch token $\mathbf{x}_p^i \in \mathbb{R}^P$ is projected into a K dimensional embedding space via linear layer $\theta \in \mathbb{R}^{P \times K}$:

$$\mathbf{f}_0 = [\mathbf{x}_p^1 \theta, \mathbf{x}_p^2 \theta, \dots, \mathbf{x}_p^{H'W'} \theta]. \quad (5)$$

Different from previous global [12,2] and local attention [24], our structural attention aggregates contexts within the main anatomy. In other words, the query scope S consists of tokens in the foreground anatomy. To realize this, we use a lightweight patch classifier to classify each token as foreground or background. The classifier uses a 1×1 convolution, followed by a Sigmoid function. The result is a binary mask $M \in \mathbb{R}^{H' \times W'}$, where each element $M_{h,w}$ represents the probability of the token at position (h, w) being foreground.

To train the patch classifier, the ground-truth binary mask \hat{M} is extracted by leveraging foundational Segment-Anything Model (SAM) [14]. As SAM is a class-agnostic segmentation, we use it to extract all masks in the image and choose the one covering the largest foreground area, indicative of the primary anatomy. We compute the ground-truth masks offline. During training, they are loaded to compute the binary cross-entropy (BCE) loss between ground-truth and predicted masks to optimize the patch classifier:

$$\mathcal{L}_{\text{mask}} = -\frac{1}{H' \times W'} \sum_{i=1}^{H'} \sum_{j=1}^{W'} \left[\hat{M}_{ij} \log(M_{ij}) + (1 - \hat{M}_{ij}) \log(1 - M_{ij}) \right]. \quad (6)$$

The decoder receives the contextual features from ST blocks and generates the synthetic image $\tilde{\mathbf{y}}$ with skip connections. Each decoder layer gradually upsamples encoded features using a sequence of 2×2 deconvolution and a 3×3 convolution.

The final generator loss combines the CycleGAN’s loss [33], including cycle consistency and adversarial losses, and the mask loss, weighted by λ :

$$\mathcal{L} = \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{cycle}} + \lambda \mathcal{L}_{\text{mask}}. \quad (7)$$

Details of our Structural Transformer block. To guide the Transformer under unpaired training, our Structural Transform applies dual attention strategy. For the foreground, we perform structural attention to learn relationships between anatomical regions. For the background, the local attention is performed,

enabling effective information exchange between foreground and background features. The ST block uses the binary mask M to classify foreground tokens given a threshold σ . For a query token q_{fg} in the foreground, our model aggregates key-value tokens within a scope S_{fg} within the foreground anatomy:

$$S_{fg} = \{(i, j) \mid M_{i,j} > \sigma\}. \quad (8)$$

A default value $\sigma = 0.5$ is used to classify foreground tokens. The structural attention consists of a layer normalization (LN) followed by a multi-head self-attention (MSA) layer [26] and a multi-layer perceptron (MLP) with skip connections to facilitate the preservation of low-level features. The foreground token output \mathbf{f}_{t+1}^{fg} from the structural attention in the t -th ST block is formulated as:

$$\mathbf{f}'_{t+1}{}^{fg} = \text{SA}(\text{LN}(\mathbf{f}_t^{fg}), S_{fg}) + \mathbf{f}_t^{fg}, \text{ and} \quad (9)$$

$$\mathbf{f}_{t+1}^{fg} = \text{MLP}(\text{LN}(\mathbf{f}'_{t+1}{}^{fg})) + \mathbf{f}'_{t+1}{}^{fg}. \quad (10)$$

For a query token q_{bg} in the background, we aggregate tokens around a local window $m \times m$:

$$S_{bg} = \{(i, j) \mid i \in [i_q - \frac{m}{2}, i_q + \frac{m}{2}], j \in [j_q - \frac{m}{2}, j_q + \frac{m}{2}]\}. \quad (11)$$

Similarly, the background token output \mathbf{f}_{t+1}^{bg} from the local attention is formulated as:

$$\mathbf{f}'_{t+1}{}^{bg} = \text{SA}(\text{LN}(\mathbf{f}_t^{bg}), S_b) + \mathbf{f}_t^{bg}, \text{ and} \quad (12)$$

$$\mathbf{f}_{t+1}^{bg} = \text{MLP}(\text{LN}(\mathbf{f}'_{t+1}{}^{bg})) + \mathbf{f}'_{t+1}{}^{bg}. \quad (13)$$

Lastly, the dual-attention ST block consolidates the output \mathbf{f}_{t+1} based on the indices (h, w) designated for the foreground S_{fg} , and the background scope S_{bg} .

3 Experimental Results

Datasets and implementation details. For MR-to-CT and MR-to-PET translation tasks, experiments are conducted on the MRXFDG [18] dataset, a public neuroimaging set containing cerebral scans from 37 adult subjects. Sagittal slices are extracted and resized to 224×224 pixels. For the PET-to-CT translation task, models are evaluated on the *AutoPET* dataset [10], comprising 310 whole-body PET and CT scans acquired from a Siemens Biograph mCT scanner. Coronal views are extracted and resized to 256×256 pixels. Both datasets are partitioned into training, validation, and test sets with the proportion of 8:1:1.

All models are trained using the Adam optimizer for 100 epochs, with a learning rate of $1e-4$ which linearly decays to zero over the last 50 epochs. We train with a batch size of 16 on two NVIDIA RTX 3090 GPUs.

Evaluation metrics. Employing standard performance metrics from prior studies [27,15], we compute mean absolute error (MAE), peak signal-to-noise ratio

(PSNR), and structural similarity (SSIM) between the real and synthesized images. The reported results are averaged over five runs. A paired student’s t -test is conducted between UNest and the compared methods to test the significance of performance difference ($p = 0.05$).

Table 1. Comparison of different methods on head scan, MRXFDG dataset [18] for four translation tasks. The improvement of UNest over all compared methods is statistically significant at $p = 0.05$.

Methods	MR-to-PET			PET-to-MR			MR-to-CT			CT-to-MR		
	MAE ↓	PSNR ↑	SSIM ↑	MAE ↓	PSNR ↑	SSIM ↑	MAE ↓	PSNR ↑	SSIM ↑	MAE ↓	PSNR ↑	SSIM ↑
<i>Convolution-based</i>												
CycleGAN [33]	7.88±.1	33.11±.01	79.53±.2	11.21±.1	31.89±.02	62.47±.3	7.20±.1	34.67±.04	82.56±.5	10.75±.2	32.15±.06	70.54±.7
sc-CycleGAN [27]	7.54±.1	33.20±.02	80.06±.3	10.64±.2	31.93±.03	62.87±.3	7.12±.2	34.82±.05	83.06±.7	10.34±.4	32.33±.11	70.77±.8
AttentionGAN [23]	7.21±.1	33.54±.03	81.09±.2	10.23±.1	33.02±.02	70.97±.1	6.84±.1	34.97±.05	83.88±.6	9.92±.4	33.12±.09	72.09±.9
MaskGAN [19]	6.95±.1	33.98±.02	81.76±.2	9.81±.1	33.10±.02	71.35±.2	6.59±.2	35.10±.05	84.53±.7	9.37±.3	33.39±.09	72.53±.1.0
<i>Hybrid Conv-Trans</i>												
TransUNet [2]	13.10±.2	32.68±.04	70.81±.6	13.70±.3	31.78±.03	59.25±.5	10.25±.2	33.47±.04	75.52±.8	12.48±.3	31.22±.06	62.32±.9
ResViT [6]	10.44±.1	32.99±.03	76.84±.5	14.66±.2	30.62±.03	56.24±.4	7.58±.2	34.72±.03	81.89±.7	11.72±.3	31.85±.07	64.51±.3
Cytran [22]	8.11±.1	33.23±.02	78.29±.5	11.05±.3	32.22±.04	60.74±.5	6.98±.1	34.89±.04	82.68±.7	9.85±.2	32.45±.06	70.80±.1.0
<i>Pure Transformers</i>												
PTNet [32]	11.43±.3	32.76±.07	70.53±.8	13.63±.4	32.33±.08	63.21±.9	9.42±.4	33.90±.10	75.88±.9	12.08±.7	31.43±.13	62.89±.1.2
UNETR [12]	7.35±.1	33.60±.04	81.90±.4	8.98±.1	33.06±.03	72.18±.2	7.95±.4	34.55±.11	81.18±.9	10.45±.3	32.98±.05	68.33±.1.0
Swin-UNETR [24]	7.54±.2	33.53±.03	81.72±.6	9.11±.2	33.03±.07	71.49±.8	8.80±.7	34.28±.18	80.53±.1.11	11.15±.5	32.87±.05	67.93±.1.3
UNest (Ours)	6.32±.2	34.24±.04	82.97±.4	8.10±.3	33.91±.05	74.13±.3	6.18±.3	35.76±.09	85.48±.5	8.72±.4	34.17±.07	72.97±.9




Fig. 4. Visual results of different methods for MRI-to-PET and MRI-to-CT translation on MRXFDG dataset.

Comparisons with state-of-the-art. Table 1 benchmarks four synthesis tasks on the MRXFDG dataset. Without *inductive bias*, UNETR is inferior to convolution AttentionGAN on a small-sized dataset, which is common in the medical domain. Injecting strong learning biases, the proposed UNest outperforms

Swin-UNETR by 19.30% and 3.95% respectively in terms of MAE and PSNR. Fig. 4 presents visual results on MR-to-PET and MR-to-CT. Without inductive bias, UNETR tends to produce blurrier details, while Swin-UNETR distorts the details in the subcortical brain structures. In contrast, our structural attention method, accurately synthesizes the structural regions in PET. Table 2 presents the benchmark of different synthesis methods on AutoPET dataset [10] for PET-to-CT and CT-to-PET tasks. Incorporating a structural inductive bias, our UNest significantly improves upon UNETR and Swin-UNETR by 14.21% and 11.73% respectively in terms of MAE.

Table 2. Quantitative comparison of different methods on the AutoPET dataset for translating between PET and CT. The improvement of UNest over all compared methods is statistically significant at $p = 0.05$. FG-S + BG-S: an ablated version of UNest with structural attention applying on both foreground and background regions.

Methods	PET-to-CT			CT-to-PET		
	MAE ↓	PSNR ↑	SSIM ↑	MAE ↓	PSNR ↑	SSIM ↑
<i>Convolution-based</i>						
CycleGAN [33]	14.34±.2	32.30±.04	66.85±.6	12.80±.1	31.95±.02	69.53±.5
sc-CycleGAN [27]	14.59±.2	32.42±.06	67.43±.8	12.77±.1	32.05±.04	70.53±.8
AttentionGAN [23]	13.62±.2	32.97±.06	67.12±.6	11.87±.1	32.61±.03	72.54±.6
QS-Attention [13]	13.21±.4	33.70±.08	71.49±1.1	9.75±.2	33.21±.07	77.42±1.0
<i>Hybrid Convolution-Transformer</i>						
TransUNet [2]	16.65±.5	32.68±.10	62.68±1.1	13.52±.4	31.94±.09	67.88±1.1
Cytran [22]	14.77±.3	32.24±.08	64.37±.9	9.05±.2	33.56±.06	74.65±.8
ResViT [6]	13.55±.3	33.28±.07	70.53±.7	9.90±.1	32.13±.06	68.29±.7
<i>Pure Transformers</i>						
PTNet [32]	14.11±.6	32.62±.10	66.53±1.3	9.24±.5	32.77±.09	70.53±1.1
UNETR [12]	11.08±.3	33.70±.08	76.06±.8	7.77±.2	33.58±.05	80.30±.8
Swin-UNETR [24]	10.93±.4	33.78±.08	76.42±.9	7.42±.3	33.78±.07	80.53±.8
UNest (Ours)	9.57±.3	34.05±.07	78.47±.8	6.55±.2	34.27±.05	81.55±.7
UNest w/ FG-S + BG-S	9.73 ±.4	33.85 ±.07	77.85 ±.8	6.87 ±.3	34.11 ±.05	81.21 ±.8

Ablation study. Table 2 shows that using dual attention with structural attention FG-S + BG-S improves MAE of UNETR and Swin UNETR by 12.18% and 10.98%, respectively, on PET-CT. Separating attentions on BG and FG injects a structural inductive bias, useful for unpaired training. *Notably*, using our hybrid attention with FG-S and BG-L achieves the best results, which is further supported by visual results in Fig. 5-Left. Fig. 5-Right shows that global attention attends to less relevant BG tokens, while our structural attention adaptively attends to anatomical features. It incorporates a long-range context for smooth brain regions (top) and a localized focus for a structural sinus (bottom).

4 Conclusion

In this paper, we introduce the UNet Structural Transformer (UNest) for unpaired medical image synthesis. Traditional Transformer methods often under-

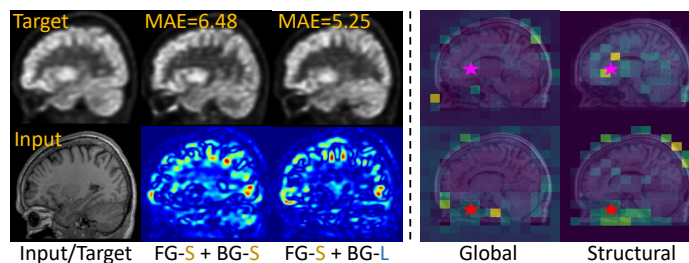


Fig. 5. *Left:* Error maps of synPET produced by UNest with FG-S + BG-S attention, and our hybrid S-L attention. *Right:* Attention maps of global and structural attention.

perform in unpaired image synthesis, due to the absence of appropriate learning priors. Injecting structural priors as in our UNest significantly reduces the synthesis error of previous Transformer models by 19.3% on the small-sized MRXFDG. The structural inductive bias also boosts synthesis performance of current models by 14.2% on the large-scale AutoPET dataset.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Learning with radiomics for disease diagnosis and treatment planning: A review. *Frontiers in Oncology* **12**, 773840 (2022)
2. Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y.: TransUNet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306* (2021)
3. Chen, X., Wang, X., Zhou, J., Qiao, Y., Dong, C.: Activating more pixels in image super-resolution transformer. In: *CVPR*. pp. 22367–22377 (2023)
4. Cui, C., Yang, H., Wang, Y., Zhao, S., Asad, Z., Coburn, L.A., Wilson, K.T., Landman, B., Huo, Y.: Deep multi-modal fusion of image and non-image data in disease diagnosis and prognosis: a review. *Prog. in Biomed. Engineer.* (2023)
5. Dai, Y., Gao, Y., Liu, F.: Transmed: Transformers advance multi-modal medical image classification. *Diagnostics* **11**(8), 1384 (2021)
6. Dalmaz, O., Yurt, M., Çukur, T.: Resvit: residual vision transformers for multi-modal medical image synthesis. *IEEE Trans. Med. Imag.* **41**(10), 2598–2614 (2022)
7. Doherty, D., Millen, K.J., Barkovich, A.J.: Midbrain and hindbrain malformations: advances in clinical diagnosis, imaging, and genetics. *The Lancet Neurology* **12**(4), 381–393 (2013)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *Proc. Int. Conf. Learn. Represent.* (2021)
9. Emami, H., Dong, M., Nejad-Davarani, S.P., Glide-Hurst, C.K.: SA-GAN: Structure-aware gan for organ-preserving synthetic ct generation. In: *MICCAI*. pp. 471–481. Springer (2021)

10. Gatidis, S., Hepp, T., Früh, M., La Fougère, C., Nikolaou, K., Pfannenber, C., Schölkopf, B., Küstner, T., Cyran, C., Rubin, D.: A whole-body fdg-pet/ct dataset with manually annotated tumor lesions. *Scientific Data* **9**(1), 601 (2022)
11. Ge, Y., Wei, D., Xue, Z., Wang, Q., Zhou, X., Zhan, Y., Liao, S.: Unpaired MR to CT synthesis with explicit structural constrained adversarial learning. In: *IEEE Int. Symposium Biomed. Imaging. IEEE* (2019)
12. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: *WACV*. pp. 574–584 (2022)
13. Hu, X., Zhou, X., Huang, Q., Shi, Z., Sun, L., Li, Q.: Qs-attn: Query-selected attention for contrastive learning in i2i translation. In: *CVPR*. pp. 18291–18300 (2022)
14. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv preprint arXiv:2304.02643* (2023)
15. Liu, Y., Chen, A., Shi, H., Huang, S., Zheng, W., Liu, Z., Zhang, Q., Yang, X.: CT synthesis from MRI using multi-cycle GAN for head-and-neck radiation therapy. *Computer. Med. Imag. Graphic.* **91**, 101953 (2021)
16. Lu, Z., Xie, H., Liu, C., Zhang, Y.: Bridging the gap between vision transformers and convolutional neural networks on small datasets. *Proc. Adv. Neural Inform. Process. Syst.* **35**, 14663–14677 (2022)
17. Matsuo, H., Nishio, M., Nogami, M., Zeng, F., Kurimoto, T., Kaushik, S., Wiesinger, F., Kono, A.K., Murakami, T.: Unsupervised-learning-based method for chest MRI-CT transformation using structure constrained unsupervised generative attention networks. *Scientific reports* **12**(1), 11090 (2022)
18. Mérida, I., Jung, J., Bouvard, S., Le Bars, D., Lancelot, S., Lavenne, F., Bouillot, C., Redouté, J., Hammers, A., Costes, N.: Cermep-idb-mrxfdg: a database of 37 normal adult human brain [18f] fdg pet, t1 and flair mri, and ct images available for research. *EJNMMI research* **11**(1), 1–10 (2021)
19. Phan, V.M.H., Liao, Z., Verjans, J.W., To, M.S.: Structure-preserving synthesis: Maskgan for unpaired mr-ct translation. In: *MICCAI*. pp. 56–65. Springer (2023)
20. Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., Dosovitskiy, A.: Do vision transformers see like convolutional neural networks? *Proc. Adv. Neural Inform. Process. Syst.* **34**, 12116–12128 (2021)
21. Richardson, D.B., Cardis, E., Daniels, R.D., Gillies, M., O’Hagan, J.A., Hamra, G.B., Haylock, R., Laurier, D., Leuraud, K., Moissonnier, M., et al.: Risk of cancer from occupational exposure to ionising radiation: retrospective cohort study of workers in france, the united kingdom, and the united states. *BMJ* **351** (2015)
22. Ristea, N.C., Miron, A.I., Savencu, O., Georgescu, M.I., Verga, N., Khan, F.S., Ionescu, R.T.: Cytran: A cycle-consistent transformer with multi-level consistency for non-contrast to contrast ct translation. *Neurocomputing* (2023)
23. Tang, H., Liu, H., Xu, D., Torr, P.H., Sebe, N.: Attentiongan: Unpaired image-to-image translation using attention-guided generative adversarial networks. *IEEE Trans. Neu. Netw. Learn. Syst.* (2021)
24. Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A.: Self-supervised pre-training of swin transformers for 3d medical image analysis. In: *CVPR* (2022)
25. Torbunov, D., Huang, Y., Yu, H., Huang, J., Yoo, S., Lin, M., Viren, B., Ren, Y.: Uvcgan: Unet vision transformer cycle-consistent gan for unpaired image-to-image translation. In: *WACV*. pp. 702–712 (2023)

26. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need **30** (2017)
27. Yang, H., Sun, J., Carass, A., Zhao, C., Lee, J., Prince, J.L., Xu, Z.: Unsupervised MR-to-CT synthesis using structure-constrained CycleGAN. *IEEE Trans. Med. Imag.* **39**(12), 4249–4261 (2020)
28. Yang, H., Sun, J., Carass, A., Zhao, C., Lee, J., Xu, Z., Prince, J.: Unpaired brain MR-to-CT synthesis using a structure-constrained CycleGAN. In: *MICCAI Workshop: Deep Learn. Med. Image Analysis*, pp. 174–182. Springer (2018)
29. Yu, F., Wang, X., Cao, M., Li, G., Shan, Y., Dong, C.: Osrt: Omnidirectional image super-resolution with distortion-aware transformer. In: *CVPR*. pp. 13283–13292 (2023)
30. Zhang, B., Liu, L., Phan, M.H., Tian, Z., Shen, C., Liu, Y.: Segvitv2: Exploring efficient and continual semantic segmentation with plain vision transformers. *Int. J. Comput. Vis.* (2023)
31. Zhang, J., Cui, Z., Jiang, C., Zhang, J., Gao, F., Shen, D.: Mapping in cycles: Dual-domain PET-CT synthesis framework with cycle-consistent constraints. In: *MICCAI*. pp. 758–767. Springer (2022)
32. Zhang, X., He, X., Guo, J., Eftehadi, N., Aw, N., Semanek, D., Posner, J., Laine, A., Wang, Y.: Ptnet3d: A 3d high-resolution longitudinal infant brain mri synthesizer based on transformers. *IEEE Trans. Med. Imag.* **41**(10), 2925–2940 (2022)
33. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *ICCV*. pp. 2223–2232 (2017)